Principal Associate Directorate for Science, Technology, and Engineering (PADSTE)

Theory, Simulation & Computation (ADTSC)
High Performance Computing (HPC) Division

A partial listing of ISTI Interactions

IS&T DATA EXPORATORY - DATA INTENSIVE SUPERCOMPUTING AT LANL

The IS&T Data Exploratory is a unique data intensive supercomputing facility designed to explore the challenges of working with large datasets. Sensors, Internet packet filters, telescopes and satellites all generate massive amounts of data. To solve problems in areas like energy security, bio-security, and cosmology we need effective ways to analyze this data. The Data Exploratory will provide a trial facility in an open environment where researchers can experiment with data intensive methods.

One of the challenges of large datasets is representing the data in a manageable format that scientists can learn from. The IS&T Data Exploratory will include a visualization environment, the DISC Visualization Collaboratory, to provide an alternative method for looking at results from data intensive applications. The goal is to find new ways of using visualization for information applications, to provide an abstraction of the data rather than a simulation of it.

ISTI COLLABORATIVE RESEARCH PROGRAM

CRP Project (LANL/CMU): ...And eat it too: High read performance in write-optimized HPC I/O middleware file formats

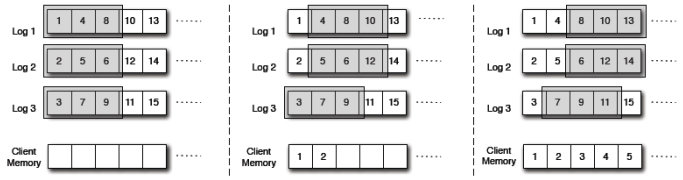

'Progress of sequential reading of a checkpoint log file incurs no additional seeks.'

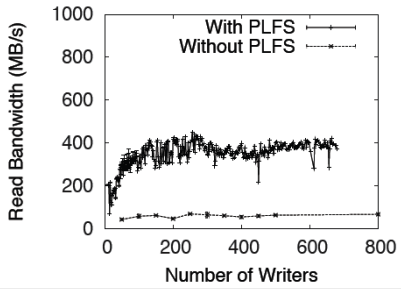

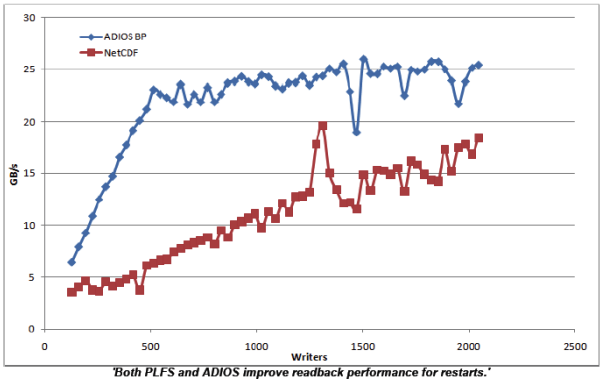

'Both PLFS and ADIOS improve readback performance for restarts.'

As HPC applications run on increasingly high process counts on larger and larger machines, both the frequency of checkpoints needed for fault tolerance and the resolution and size of Data Analysis Dumps are expected to increase proportionally. In order to maintain an acceptable ratio of time spent performing useful computation work to time spent performing I/O, write bandwidth to the underlying storage system must increase proportionally to this increase in the checkpoint and computation size. Unfortunately, popular scientific self-describing file formats such as netCDF and HDF5 are designed with a focus on portability and exibility. Extra care and careful crafting of the output structure and API calls is required to optimize for write performance using these APIs. To provide sufficient write bandwidth to continue to support the demands of scientific applications, the HPC community has developed a number of I/O middleware layers, that structure output into write-optimized file formats. However, the obvious concern with any write optimized file format would be a corresponding penalty on reads. In the log-structured filesystem, for example, a file generated by random writes could be written efficiently, but reading the file back sequentially later would result in very poor performance. Simulation results require efficient read-back for visualization and analytics, and though most checkpoint files are never used, the efficiency of a restart is very important in the face of inevitable failures. The utility of write speed improving middleware would be greatly diminished if it sacrificed acceptable read performance. In this paper we examine the read performance of two write-optimized middleware layers on large parallel machines and compare it to reading data natively in popular file formats.

LANL (HPC) Collaborators/mentors: James Nuñez, HPC-5; John Bent, HPC-5; Meghan Wingate, HPC-5; Ben McClelland, HPC-5

Related Papers/presentations

- Paper: SC2009:...And eat it too: High read performance in write-optimized HPC I/O middleware file formats, November 2009
- Poster: USENIX FAST 10: Enabling Scientific Application I/O on Cloud FileSystems, February 2010
- Presentation: Fast Log-based Concurrent Writing of Checkpoints PDSW 08, May 2008
- Poster: Fast Log-based Concurrent Writing of Checkpoints PDSW 08, November 2008
- Presentation: Log-structured Files for Fast Checkpointing PDSW08, October 2008
- Poster: Log-structured Files for Fast Checkpointing PDSW08, October 2008
- Poster: USENIX FAST 10—Enabling Scientific Application I/O on Cloud FileSystems, February 2010

CRP Project (LANL/UCSC): Alternative Reliability Models in Ceph

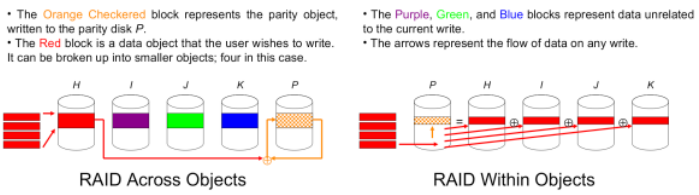

RAID systems have traditionally offered increased performance and data security in small storage systems. An opportunity now exists to extend traditional RAID principles into the area of large-scale object-based storage devices in order to offer greater data security and space efficiency. In a system where component failures can be expected on a daily basis, the importance of redundancy mechanisms is obvious, and RAID systems and to learn how they function in a new environment.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5, HB Chen, HPC-5

Related Papers/presentations

- Paper: Mahanaxar: Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010
- Poster: SSRC07 Alternative Reliability in Ceph, 2007
- Presentation: ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2009
- Presentation: 26th IEEE Symposium (MSST2010): Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010

CRP Project (LANL/UCSC): COLT: Continuous On-Line Tuning of Databases

Tuning is one of the most challenging and important tasks in setting up a database system. A typical assumption is that a representative workload can provide the means to perform some initial tuning automatically, which can then be refined manually by the database administrator. However, several emerging application domains challenge this assumption. A characteristic example is scientific data management, where the use of ad-hoc data analytics cancels the ability to gather a representative workload, and the lack of database expertise among scientists lowers the possibility for efficient manual tuning. This motivates us to look into autonomic tuning techniques that operate continuously and require minimal intervention from an administrator. In particular, we examine online methods to tune the physical database design, i.e., the set of materialized data structures, such as indices and materialized views that are crucial for efficient query processing.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- Paper: PostgreSQL (To appear as a chapter in Database System Concepts, 6th ed.), 2009
- Paper: On-line Index Selection for Shifting Workloads SMDB 2007, April 2007
- Paper: Index Interactions in Physical Design Tuning: Modeling, Analysis, and Applications, August 2009
- Paper: On-line Index Selection for Shifting Workloads SMDB 2006, 2006
- Paper: COLT: Continuous On-Line Tuning, 2006
- Paper: SIGMOD/PODS 2010—An Automated, yet Interactive and Portable DB designer, June 2010
- Paper: Semi-Automatic Index Tuning: Keeping DBAs in the Loop, 2010
- Presentation: ISSDM Day: Physical Database Tuning with Interaction-Aware Index Selection, October 2009
- Presentation: SIGMOD 2010—An Automated, yet Interactive and Portable DB designer, June 2010
- Presentation: 2010 ISSDM Day: On-line Index Selection for Physical Database Tuning, October 2010
- Poster: 2010 ISSDM Day: On-line Index Selection for Physical Database Tuning, October 2010

CRP Project (LANL/UCSC): Concurrent Programming Models for the Cell Broadband Engine

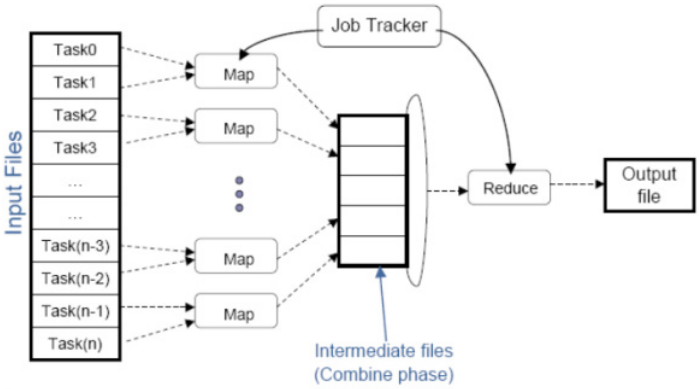

The Cell is powerful, but challenging to program. Concurrent Programming Models for the Cell Broadband Engine may help. Everything old is new again, Cilk can provide a smooth transition of older C codes to the Cell.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO

CRP Project (LANL/CMU): Data Intensive Super Computing for Science

This project explores how well suited data intensive computing programming/run time paradigms like map reduce and other graphs apply to scientific applications.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5; Tony Heaton, HPC-5; Laura Monroe, HPC-5

Related Papers/presentations

- Paper: Introducing Map-Reduce to High End Computing PDSW08, November 2008
- Paper: DiskReduce: RAID for Data-Intensive Scalable Computing, November 2009
- Paper: In Search of an API for Scalable File Systems: Under the Table or Above It?, June 2009

CRP Project (LANL/UCSC): Dynamic Load-Balancing in Petabyte-scale File Systems that use Pseudo-Random Data

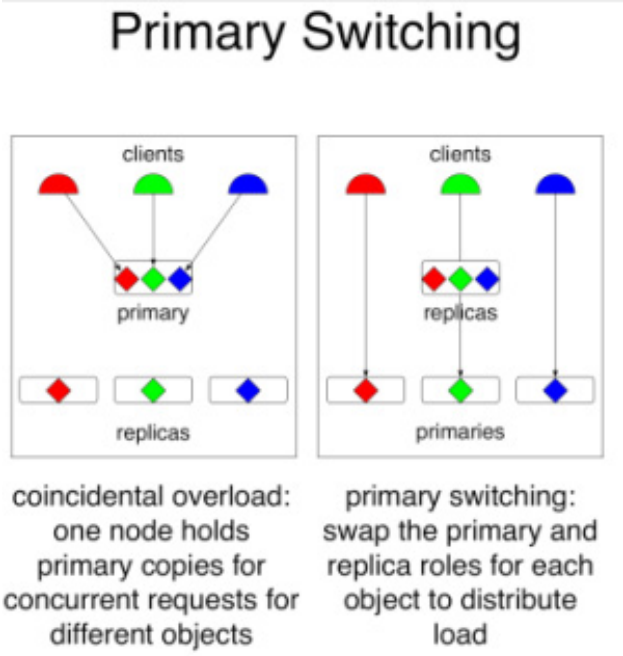

Pseudorandom placement in distributed storage systems offers scalability benefits. Pseudorandom placement makes load balancing harder; new techniques are required. We explore different load balancing techniques using Ceph, an object-based storage system developed at UCSC.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; Alfred Torrez, HPC-1; John Bent, HPC-5; James Nuñez, HPC-5; Meghan Wingate, HPC-5

Related Papers/presentations

- Paper—SciDAC 2009: Building a parallel file system simulator, June 2009
- Presentation—USENIX FAST08 WIP: Dynamic Load Balancing in Ceph, February 2008
- Presentation—ISSDM Day: Performance Analysis of Mixed Distributed Filesystem Workloads, October 2009
- Paper: Building a parallel file system simulator
- Presentation: SSRC08 Scaling in Ceph, May 2008
- Presentation: SSRC07 Hadoop on Ceph, 2007
- Presentation; PDL09:PLFS and HDFS: Enabling Parallel Filesystems Semantics in the Cloud, November 2009
- Presentation: 12010 Mini-Showcase—Scalable Simulation of Parallel Filesystems, August 2010

CRP Project (LANL/UCSC): Erasure Codes

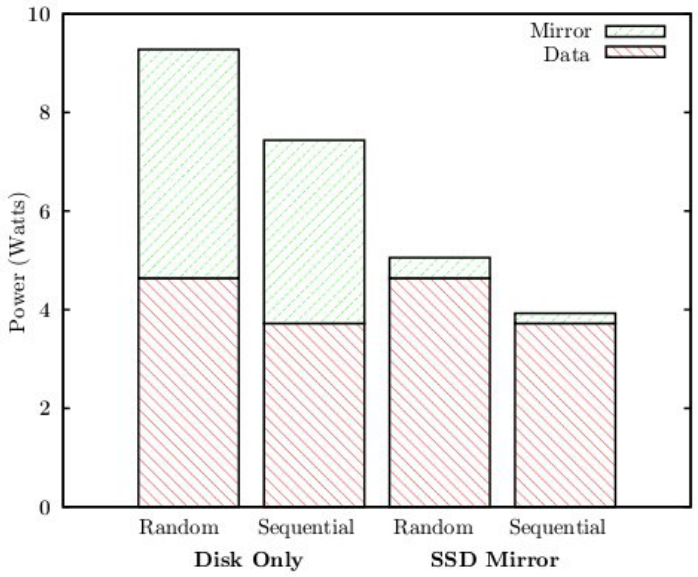

Large data centers are typically composed of many hard drives arranged in a striped layout with redundancy, to provide adequate performance and reliability. Solid state disk drives (SSDs) are a newer technology that has better performance and lower power requirements compared to hard drives. This research examines the use of solid state disk drives in disk-SSD hybrid storage systems.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO, James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- Paper: Self-Adaptive Two-Dimensional RAID Arrays. IPCCC 2007:246-253, 2007
- Paper: Disaster recovery codes: increasing reliability with large-stripe erasure correcting codes. StorageSS 2007:31-36, 2007
- Paper: HotEMNets Energy-Reliability Trade-offs in Sensor Networks, June 2008
- Poster: UCSC Engr Day 2009: Reducing Power Consumption by Incorporating Flash Memory into RAID Storage, May 2009
- Presentation: ISSDM Day: Improving RAID-Based Storage Systems with Flash Memory, October 2009
- Poster: EuroSys 2010—RAID4S: Adding SSDs to RAID Arrays, April 2010
- Poster: USENIX FAST 10—RAID4S: Adding SSDs to RAID Arrays, February 2010

CRP Project (LANL/UCSC): Exascale Data Management in File Systems

Today's file systems use interfaces largely conforming to POSIX IO, a standard that was designed in the mid-1960s when high-end file systems stored less than 100MB. Today's high-end file systems in super-computing environments and search engine and social network companies are up to 7-9 orders of magnitude larger, resulting in numbers of data items for which POSIX abstractions are quite inadequate. With the advent of exascale systems this inadequacy will continue to grow worse. We propose to coalesce the functionality of data analysis, management, and file systems by extending file systems with data management services, including declarative querying, distributed query planning and optimization, automatic indexing, a common data model for most scientific file formats, and provenance tracking that spans multiple layers of abstraction. These services are optimized for predictability, throughput, and low latency while minimizing data movement and power consumption.

We will leverage important insights gained by both the file system and the database communities. However, the sheer amount of data managed by file systems and their associated access patterns require a design very different from current database management systems.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; John Bent, HPC-5; Carolyn Connor, HPC-5; Lartchesar Ionkov, HPC-5; Michael Lang, HPC-5

Related Papers/presentations

- Paper: Comparing the Performance of Different Parallel Filesystem Placement Strategies, February 2009
- Paper: Abstract Storage: Moving File Format - Specific Abstractions into Petabyte-Scale Storage Systems, June 2009
- Paper: A Flexible Scheduling Framework Supporting Multiple Programming Models with Arbitrary Semantics in Linux, September 2009

CRP Project (LANL/CMU): File System and File Structure

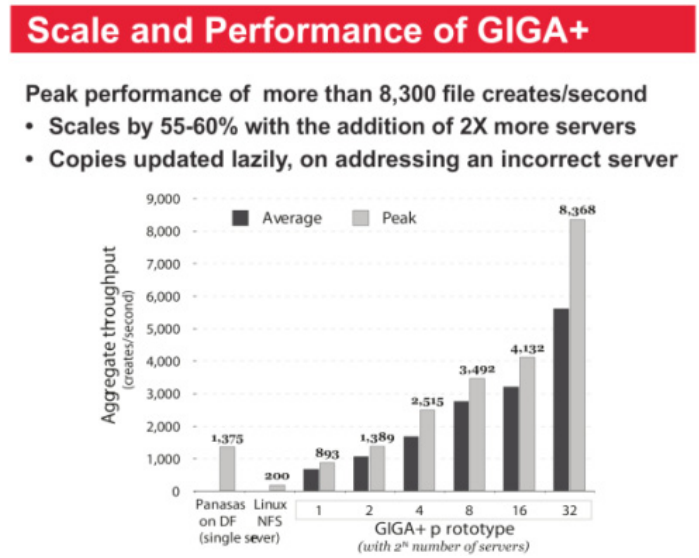

The concept is to head HPC file system storage towards file formats in the file system. It is quite possible that many file types like N processes to 1 file with small strided writes might be served well by special handling at the file system level. Decades ago, file types and access methods were used and were supported within a single file system. The IBM MVS storage systems allowed for many different file types, partitioned data sets, indexed sequential, virtual sequential, and sequential to name a few. Storage for modern HPC systems may benefit from a new parallel/scalable version of file types. There is much research to be done in this area to determine the usefulness of this concept and how such a thing would work with modern supercomputers and future HPC languages and operating environments.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; John Bent, HPC-5; James Nuñez, HPC-5

Related Papers/presentations

- Paper: SC09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- Paper: Understanding disk failure rates: What does an MTTF of 1, 000, 000 hours mean to you?. TOS 3(3) (2007), 2007
- Paper: MultiMap: Preserving disk locality for multidimensional datasets. ICDE 2007:926-935, 2007
- Paper: Verifying distributed erasure-coded data. PODC 2007:139-146, 2007
- Paper: Modeling the relative fitness of storage. SIGMETRICS 2007:37-48, 2007
- Paper: Low-overhead byzantine fault-tolerant storage. SOSP 2007:73-86, 2007
- Paper: Using Provenance to Aid in Personal File Search. USENIX Annual Technical Conference 2007:171-184, 2007
- Paper: Fast log-based concurrent writing of checkpoints PDSW08, November 2008
- Paper: WISH at ASPLOS XIV Enabling Enterprise Solid State Disks Performance" Milo Polte, Jiri Simsa, Garth Gibson. Accepted to the First Workshop on Integrating Solid-State Memory into the Storage Hierarchy, March 2009
- Paper: USENIX HotCloud09 In search of an API for scalable file systems: Under the table or above it?", June 2009tions, November 2009

- ◊ Poster: Efficient Data Placement in a Hybrid Storage Architecture PDL08, October 2008
- ◊ Poster: CMU PDL Day PLFS: A Checkpoint Filesystem for Parallel Applications, May 2009
- ◊ Presentation: CMU PDL Day Understanding Performance in Solid State Disks, May 2009
- ◊ Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
- ◊ Presentation: PDL09: Tables and File Systems: Moving Into the Cloud, November 2009
- ◊ Presentation: PDL09: GIGA+: Scalable Directories from Shared File Systems, November 2009
- ◊ Presentation: PDL09: Large-scale Evaluation of GIGA+ Scalable Directories, November 2009
- ◊ Presentation; PDL09: Reliability Modeling for Large Scale Declustered Storage, November 2009
- ◊ Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- ◊ Presentation: PDL WS 2010: PLFS on Jaguar/Lustre at Scale, October 2010

CRP Project (LANL/UCSC): Incorporating solid state drives into distributed storage systems

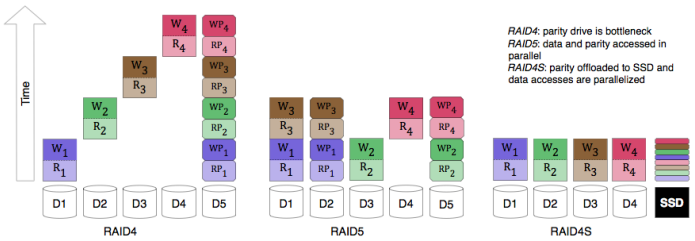

This project examines a promising architecture that uses a limited number of SSDs (solid state drives) to decrease the power consumption and either increase the performance or reliablitility of RAID storage subsystems. We are researching the use of SSDs for parity storage in a disk-SSD hybrid RAID system.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; John Bent, HPC-5; James Nuñez, HPC-5; Meghan Wingate, HPC-5

Related Papers/presentations

- ◊ Presentation: SSRC08 Self Managed Raid, 2007
- ◊ Presentation: EuroSys 2010—RAID4S: Adding SSDs to RAID Arrays, April 2010
- ◊ Presentation: 2010 Mini-Showcase—Improving RAID Performance with Solid State Drives, August 2010
- ◊ Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010

CRP Project (LANL/UCSC): Information Trust

Figure 2: Text coloring of the Italian cuisine article resulting after an anonymous author (unrelated to the PI) added a comparison of the virtues of the Italian and French cuisines (revision id: 104258868).

To help users make sense of collaboratively-generated information, we are developing algorithmic notions of information trust.

Affiliated Project: Measuring Wikipedia—Google funded.

LANL (HPC) Collaborators/mentors: Shelly Spearing, HPC-1; Jorge Roman, HPC-1

Related Papers/presentations

- ◊ Presentation: ISSDM Day: SCRAWL: A Semantic Crawler for the Wikipedia and Beyond, October 2009
- ◊ Presentation: 2010 ISSDM Day: WikiTrust Turning Wikipedia Quantity into Quality, October 2010
- ◊ Paper: ECRTS 2010: DP-FAIR: A Simple Model for Understanding Optimal Multiprocessor Scheduling, July 2010
- ◊ Paper: The Scientist: Redesigning Scientific Reputation, September 2010
- ◊ Paper: Measuring Contributions to Email-Based Discussion Groups, November 2009
- ◊ Paper: PAN 2010: Detecting Wikipedia Vandalism using WikiTrust, September 2010
- ◊ Presentation; ISSDM Day: SCRAWL: A Semantic Crawler for the Wikipedia and Beyond, October 2009
- ◊ Presentation: ISSDM Day: WikiTrust: Experience is what you get..., October 2009
- ◊ Presentation: 2010 ISSDM Day: WikiTrust Turning Wikipedia Quantity into Quality, October 2009
- ◊ Presentation: PAN 2010: Detecting Wikipedia Vandalism using WikiTrust, September 2010

CRP Project (LANL/CMU): Large Scale Parallel Tools: OpenSpeedShop Frameworks Plugins

Computer scientists at LANL have become intimately involved in studying the software development challenges posed by the new Roadrunner cell hybrid architecture, as well as possible approaches to addressing these. One of their findings is that the core structure of existing applications has to be revisited and understood before they can be effectively transitioned to a new architecture. Understanding this structure requires analysis tools to assess time, IO, memory usage, and other operational characteristics of the application so that its data, functions and memory patterns can be mapped to appropriate segments of the new processor architecture.

The purpose of this project is to study the OpenSpeedShop tool framework and identify a set of constraints it imposes on tool plugins, drawing on experience from OpenSpeedShop developers at LANL. The goal is to understand the rationale behind these constraints in terms of framework quality attributes like performance and extensibility, and to explore potential design alternatives and their tradeoffs with respect to those quality attributes.

LANL (HPC) Collaborators/mentors: David R. Montoya, HPC-3; Steve Painter, HPC-1; Scott Matthews, HPC-1

Related Papers/presentations

- ◊ Paper—ASE 2008: Error Reporting Logic, June 2008
- ◊ Paper—ECOOP 2009: Checking Framework Interactions with Relationships, July 2009
- ◊ Paper—RAOOL 2009: Retrieving Relationships from Declarative Files, July 2009
- ◊ Presentation—SPLASH/OOPSLA 2010: Verifying Configuration Files, October 2010
- ◊ Presentation: Checking Framework Interactions with Relationships, July 2009
- ◊ Presentation: Proper Plugin Protocols, February 2010
- ◊ Presentation: SPLASH/OOPSLA 2010: Verifying Configuration Files, October 2010
- ◊ Related Proposals
- ◊ Proposal: DOE SC Peta Tools: Petascale Memory Tools, October 2008

CRP Project (LANL/UCSC): Managing High-Bandwidth Real-Time Data Storage

This project was first conceived as a storage system for the Long Wavelength Array (LWA) radio telescope, and subsequently extended to incorporate several other problems with similar themes. The problem is characterized by ultra-high bandwidth data streams in which most data is "useless," but occasionally is "useful" but not recognized as such until a later period. These traits combine to form an odd set of requirements: there is far too much data to retain in the long-term, but which needs to be stored in the short term in order to determine whether or not a portion of it is useful.

In the case of the LWA, a practical example of this data pattern is that a significant astronomical event is detected at time T. In order to understand as much about the event as possible, scientists might declare that they need all of the data in the ten minutes prior to time T, as well as the data for some amount of time after time T. Similarly, in a cybersecurity example, an intrusion may be detected at time T coming from source A. In order to track the intrusion as closely as possible, network security personnel might declare that they need all IP packets coming from source A prior to time T, as far back as possible. These two examples involve very similar data patterns, but the data itself does not share similar organization. The main difference is in how the data is organized. In one case, a system must save and track large data elements without needing to search on anything more complicated than time. In the other, a system much save and track very small data elements which are highly structured, and needs to search on many aspects of each element, for example: origin, destination, and size. The disparity between these two types of data makes a general-purpose solution much more tricky.

The primary focus in this research is centered around providing and maintaining quality of service guarantees for real-time data on standard disk drives. This problem is complicated by the addition of other requirements, which must be integrated into the quality of service guarantees. Other processes must be able to access the drive, reliability mechanisms must be implemented to ensure data resiliency, the data must be indexed and quickly searchable for when preservation commands are given, and the system must be able to scale upwards to meet increasing bandwidth requirements. None of these extra requirements can be allowed to disrupt the central need for quality of service guarantees on the real-time data stream. This project also has research potential for a number of related secondary goals. Due to the nature of the data collection, an opportunity exists to test how well hard drives perform under unrelenting high- bandwidth writing over months and years. Solid state storage devices are unsuitable for managing this type of high-bandwidth data due to limited write cycles, but may be extremely useful when applied to specific indexing needs only. Non-standard error-correction codes not normally used for standard storage systems may find a new place in a system where every single write is a "large write" and reading, even for repair work, is extremely rare.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5; HB Chen, HPC-5; Meghan Wingate, HPC-5

Related Papers/presentations

- ◊ Presentation: USENIX FAST08 WIP Adapting RAID Methods for Use in Object Storage System, February 2008
- ◊ Presentation: SSRC08 Object Raid, May 2008
- ◊ Presentation: SSRC07 Alternative Reliability in Ceph, 2007
- ◊ Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Managing High-Bandwidth Real-Time Streaming Data, August 2009
- ◊ Presentation: ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2009
- ◊ Presentation: 2010 Mini-Showcase—Managing High-Bandwidth Real-Time Streaming Data, August 2010

CRP Project (LANL/UCSC): Measuring Wikipedia

Google funded project on information trust.

LANL (HPC) Collaborators/mentors: Shelly Spearing, HPC-1; Jorge Roman, HPC-1

Related Papers/presentations

- ◊ Presentation: WikiSym 2009: Measuring Wikipedia, October 2009

CRP Project (LANL/CMU): Multi-Dimensional Metadata and Massive Directories in Parallel File Systems

Large-scale storage systems often store massive data sets that are highly dimensioned and massive numbers of files in directory structures. This project explores solutions to dealing with indexing and managing these data sets.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- ◊ Presentation: PDL07 Metadata-indexed Cluster Filesystems, 2007
- ◊ Presentation: SC07-PDSW Giga+ Large Scale Directories, 2007

CRP Project (LANL/UCSC): Network Quality of Service

The RAD model separates scheduling into Resource Allocation and Dispatching.
Proven correct for CPU scheduling, RADoN applies the model to the network resource.

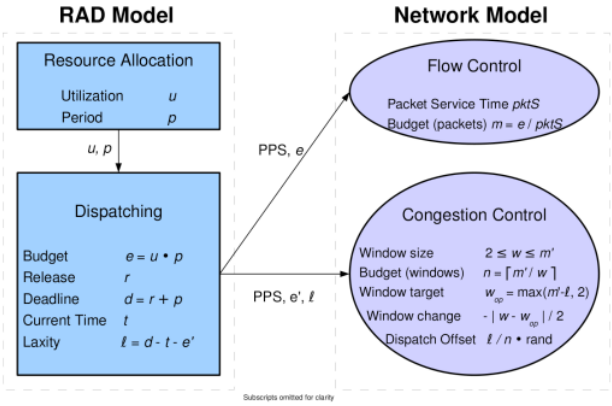

In large tightly coupled parallel systems, computation goes as fast as the slowest part. For this reason it is necessary to pursue deterministic behavior of all parts of the system. Quality of Service is one way to assist in providing deterministic behavior. This project will explore providing Quality of Service on networks of interest to high performance computing.

LANL (HPC) Collaborators/mentors: Josip Loncaric, HPC-5; Nathan Debardeleben, HPC-1

Related Papers/presentations

- ◊ Presentation: SSRC08 TCP Collapse, May 2008
- ◊ Presentation: ISSDM Day: Investigating Efficient Real-time Performance Guarantees on Storage Networks, October 2009
- ◊ Presentation: 2010 ISSDM Day: Investigating Efficient Real-Time Performance Guarantees on Storage Networks, October 2010

CRP Project (LANL/UCSC): Providing Quality of Service Support in Object-Based File Systems

In large-scale storage and file systems, multiple parallel applications can be asking for service simultaneously. In parallel workloads it is vital that deterministic service be given as the application only proceeds as fast as the slowest process. With mixed workloads this is very difficult to do in a parallel setting. We will explore how QoS support could be added to portions of the I/O stack in object based parallel file systems to enable determinism for multiple parallel applications in a mixed parallel workload environment.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- ◊ Paper: Efficient guaranteed disk request scheduling with fahrrad EuroSys 2008: 13-25, March 2008
- ◊ Paper: A Hybrid Disk-Aware Spin-Down Algorithm with I/O Subsystem Support IPCCC 2007: 236-245, 2007
- ◊ Paper: Ensuring Performance in Activity-Based File Relocation. IPCCC 2007, 2007
- ◊ Paper: Providing Quality of Service Support in Object-Based File System. MSST 2007: 157-170, 2007
- ◊ Paper: Flushing Policies for NVCache Enabled Hard Disks. MSST 2007: 299-304, 2007

- ◊ Paper: Using Comprehensive Analysis for Performance Debugging in Distributed Storage Systems. MSST 2007:281-286, 2007
- ◊ Paper: RTAS08 Virtualizing Disk Performance, July 2008
- ◊ Presentation: IEEE MSST07 Providing Quality of Service Support in Object-Based File System, 2007

CRP Project (LANL/UCSC): Scalable Security for Petascale, High Performance Storage

Storage systems are getting bigger with more high performance applications needed to access data.Data is becoming more and more confidential over time and currently security in scalable storage systems is not scalable.Systems with tens of thousands of nodes will be accessing highly distributed data using demanding I/O patterns and will need to do so securely.We intend to explore scalable security through capabilities, delegations, replication in the file system to fundamentally address scaling of security in file systems.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5

Related Papers/presentations

- ◊ Paper: Mixing Hadoop and HPC Workloads on Parallel Filesystems, November 2009
- ◊ Paper: Scalable security for petascale parallel file systems. SC 2007:16, 2007
- ◊ Paper: POTSHARDS: Secure Long-Term Storage Without Encryption. USENIX Annual Technical Conference 2007:142-156, 2008=7

CRP Project (LANL/UCSC): Scalable Simulation of Parallel File Systems

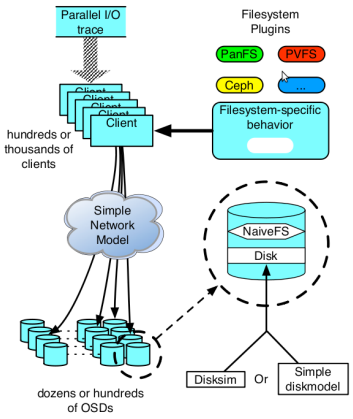

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5; HB Chen, HPC-5; Meghan Wingate, HPC-5

Related Papers/presentations

- ◊ Paper: Mixing Hadoop and HPC Workloads on Parallel Filesystems, November 2009
- ◊

CRP Project (LANL/UCSC): Search-Oriented Interfaces for Distributed Metadata Indices

Data users need effective means to locate their files in huge and increasing scaled out file systems containing millions or even billions of files. Use of simple directories is not sufficient for managing this number of files. Users already have search tools for file systems but they are based on directory structure oriented searches. We propose to explore how indexing of metadata about files could provide users more utility for searching for files in extremely large parallel file systems. We will concentrate on the interfaces users would use to search indexed metadata about files.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- ◊ Paper: SSRC Tech Report Graffiti Server - Design and Implementation, 2007

CRP Project (LANL/UCSC): Synthetic Parallel Applications

Storage system design and optimization require usage data in order to determine where in the system to focus resources. Typically this usage data is created by observing the response of the system while running synthetic benchmarks designed to place high stress on particular aspects of the system. The problem with many benchmarks is that they encourage improvements to areas of the system that do not improve actual user performance on typical workloads. Ideally, the benchmarks should be similar to actual applications so improvements can much more easily be made to storage systems. User applications often are proprietary or secure so using user applications for benchmarks is often not practical. This project will explore how synthetic benchmark applications can be synthetically generated using traces of real parallel user applications.

LANL (HPC) Collaborators/mentors: Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5

Related Papers/presentations

- ◊ Presentation: SSRC07 Hidden Markov Modeling of Parallel I/O, 2007
- ◊ Presentation: Energy-Reliability Tradeoffs in Sensor Network Storage HOTEMNETS08, June 2008

CRP Project (LANL/CMU): TCP/IP Networking Incast

LANL (HPC) Collaborators/mentors: HB Chen, HPC-5; Andrew Shewmaker, HPC-5; Parks Fields, HPC-5

RESEARCH PROPOSALS

- ◊ DOE Office of Science SciDAC2 Petascale Data Storage Institute, 2007 (Gary Grider; HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- ◊ DOE Office of Science FASTOS IO Scalable Forwarding Layer, 2007 (Gary Grider; HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- ◊ Multi-Dimensional Information Indexing Research for Observational Science Applications for Air Force, March 2008 (Gary Grider; HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- ◊ NSF with NMT: Quality of Security, 2007 (Gary Grider; HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- ◊ NSF: PRObE – The NSF Parallel Reconfigurable Observational Environment for Data Intensive Super-Computing and High Performance Computing, December 2008 (Gary Grider, HPC-DO; Carolyn Connor, HPC-5; Tommy Stup, HPC-3)
- ◊ NSF: Pathways for Undergraduate Education (CPATH) Parallel Programming, April, 2009 (Gary Grider, HPC-DO; Carolyn Connor, HPC-5)
- ◊ NSF PRObE Systems Computer Science Research Center, March 2009 (Gary Grider, HPC-DO)
- ◊ DOE SciDAC PDSI PLFS Plus Up, April 2009 (Gary Grider; HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- ◊ NSF I/O Tracing Replay - Colorado School of Mines Summer Research Program, April 2009 (Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5; Meghan Wingate, HPC-5)
- ◊ DOE FaST - Trace Replay and Visual Tools, April 2009 Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5; Meghan Wingate, HPC-5)
- ◊ LDRD: Data to Knowledge at Scale: The Power of Data-Intensive Super-Computing (LDRD-DR Preproposal #20100029DR), January 2009 (James Nuñez, HPC-5; John Bent, HPC-5; Meghan Wingate, HPC-5)
- ◊ LDRD: Data to Knowledge at Scale: The Power of Data-Intensive Super-Computing (LDRD-DR Full Proposal #20100029), March 2009 (James Nuñez, HPC-5; John Bent, HPC-5; Meghan Wingate, HPC-5)
- ◊ DOE ASC: Co-Design for Advanced Architectures and Critical Technologies Toward Exascale Computing, April 2010 (David Dubois, HPC-5; Carolyn Connor, HPC-5; Andy Dubois, HPC-5)
- ◊ LDRD Reserve Proposal: Building the Astroinformatics Competency: Finding, Interrogating, and Understanding Cosmic Explosions, August 2010 (David Dubois, HPC-5; Andy Dubois, HPC-5)
- ◊ LDRD-DR: Building the Astroinformatics Competency: Finding, Interrogating, and Understanding Cosmic Explosions, February 2010 (David Dubois, HPD-5; Andy Dubois, HPC-5)

STAFFING RECRUITING AND PIPELINE DEVELOPMENT

The Computer System, Cluster, and Networking Summer Institute (SI) is a 9-week, hands-on summer intensive program targeting Junior-level CS/CE/EE undergraduate students. SI Graduates who have worked or who are working in ADTSC include the following.

- Colby Boyer (SI2007; HPC-5)
- Ekaterina Davydenko (SI2007; currently HPC-1)
- Dane Gardner (SI 2009; HPC-3)
- Daniel Illescas (SI 2010; HPC-3)
- Jharrod LaFon (SI 2009; HPC-3)
- Evan Leeseberg (SI2008; HPC-5)
- Ruben L. Salazar (SI2008; DCS-1; currently in HPC-5)
- Lucia Short (SI2010; HPC-3)
- Graham Van Heule (SI2009; HPC-3)

COMPUTING AND INFORMATION TECHNOLOGY STUDENT MINI SHOWCASE

In its second year, the annual showcase provides a unique opportunity for students in the computing and information technology fields to present their research. The intent is to broaden students' expertise and prepare them for careers in their fields. This is an excellent forum for staff to interact and to learn about the breadth of projects in the area of IS&T at the Laboratory and for students to network and make professional contacts. A subset of CCS 2010 participants is listed below. The full listing of participants and abstracts may be found at http://institute.lanl.gov/isti/summer-school/mini-showcase/.

- **Title: Managing High-Bandwidth Real-Time Streaming Data**
  David Bigelow, University of California-Santa Cruz
  Mentors: HB Chen (HPC-5), John Bent (HPC-5)
- **Title: I/O Forwarding Scalable Layer**
  Rico D'Amore, New Mexico Highlands University
  Mentor: James Nuñez (HPC-5)
- **Title: Visualizing InfiniBand Fabric Networks with Circos**
  Karen Ho, University of California, San Diego
  Mentor: David Montoya (HPC-3)
- **Title: Memory Performance Analysis for LANL's Next Generation High Performance Computing Systems**
  Evan Leeseberg, Michigan Technological University
  Mentors: Mike Boorman (HPC-5), Andy Dubois (HPC-5), David DuBois (HPC-5)
- **Title: Compute and Data-intensive Computing: Bridging the Gap**
  Grant Mackey, University of Central Florida
  Mentor: John Bent (HPC-5)
- **Title: Data Intensive Visualization at Scale**
  Christopher Mitchell, University of Central Florida
  Mentors: James Ahrens (CCS-7), Meghan Wingate (HPC-5)
- **Title: Scalable Simulation of Parallel Filesystems**
  Esteban Molina-Estolano, University of California-Santa Cruz
  Mentor: John Bent (HPC-5)
- **Title: Infrastructure Efficiency and Environmental Monitoring**
  Alynna C. Montoya-Wiuff, New Mexico State University
  Mentor: Richard Rivera (HPC-2)
- **Title: Access Method Support File System**
  Joe Naps, University of Phoenix
  Mentors: James Nuñez (HPC-5), Meghan Wingate (HPC-5)
- **Title: Facilitating Open Science with the DISC Visualization Collaboratory**
  Katherine Nystrom, University of New Mexico
  Mentor: Carolyn Connor (HPC-5/INST-OFF)
- **Title: Upgrading the IMMS Access Grid System**
  Jake Poston, New Mexico Tech
  Mentor: Carolyn Connor (HPC-5/INST-OFF), Cindy Sievers (CCS-7)
- **Title: Detecting Silent Data Corruption**
  Ruben Salazar, University of New Mexico
  Mentors: David DuBois (HPC-5), Andrew DuBois (HPC-5)
- **Title: Extending Gazebo to Learn System Behavior for Acceptance Testing**
  Evan Samanas, University of Wisconsin
  Mentor: Craig Idler (HPC-5)
- **Title: Scalable Parallel Short-range Molecular dynamics (SPaSM)**
  David Shrader, University of Wisconsin-Madison
  Mentor: Meghan Wingate (HPC-5)
- **Title: Improving RAID Performance with Solid State Drives**
  Rosie Wacha, University of California-Santa Cruz
  Mentor: Meghan Wingate (HPC-5)
- **Title: File System Trace and Replay Redux**
  Noah Watkins, University of California-Santa Cruz
  Mentor: James Nuñez (HPC-5)
- **Title: On-Demand Scientific Computing**
  Benjamin Cheng, University of Illinois at Urbana; Gregorio Hinojos, New Mexico State University; Travis Johnson, California State University, San Bernardino
  Mentors: Andrew Shewmaker (HPC-5), Dan Gardner (HPC-3), Grahan Van Heule (HPC-3
- **Title: Performance Analysis and Evaluation of LANLs PaScalBB IO nodes using PC1e Gen 2.0 QuadData-Rate Infiniband and Multiple 10-Gigabit Ethernets**
  Juan Franco, University of N. Texas; Daniel Illescas, University of New Mexico; Rocio Perez-Medina, Northern New Mexico College
  Mentors: Alfred Torrez (HPC-1), HB Chen (HPC-5), Parks Fields (HPC-5)
- **Title: Implementation & Comparison of RDMA Over Ethernet**
  Lee Gaiser, Michigan Technological University; Brian Kraus, University of New Mexico; James Wernicke, New Mexico Mining Institute & Technology
  Mentors: Susan Coulter (HPC-3), Jharrod LaFon (HPC-3), Ben McClelland (HPC-3)
- **Title: The New Storage.lanl.gov**
  Michael Jacobi, University of Montana; Angelica Gallegos, University of New Mexico
  Mentors: Aaron Torres (HPC-3), Cody Scott (HPC-3)
- **Title: Virtualized High Performance Computing on Amazon's EC2**
  Jharrod LaFon, New Mexico State University; Dane Gardner, Golden West University
  Mentor: Dave Montoya (HPC-5)
- **Title: Data Discovery**
  Martin Loncaric, Los Alamos High School; Peter Song, Los Alamos High School

STAFF DEVELOPMENT AND RETENTION AND ORGANIZATIONAL OUTREACH ACTIVITIES

Cluster and Network Summer Institute Guest Lectures:

- Susan Coulter (HPC-3)
  - SI2008 Guest Lecture:
  - SI2009 Guest Lecture: Cluster Management
  - SI2010 Guest Lecture: IB From the Ground Up
- Cindy Martin (HPC-3)
  - SI2009 Guest Lecture: Zenoss
  - SI 2010 Guest Lecture: Zenoss
- Jharrod LaFon (HPC-3)
  - SI 2010 Guest Lecture: IB From the Ground Up
- Graham Van Heule (HPC-3)
  - SI 2010 Guest Lecture: DRM
- Cody Scott, HPC-3
  - SI 2010 Guest Lecture: Archival Storage at LANL
- John Bent, HPC-5
  - SI 2008 Guest Lecture: Parallel I/O
  - SI 2009 Guest Lecture: Parallel I/O
  - SI 2010 Guest Lecture: Parallel I/O

- HB Chen HPC-5
  - SI 2008 Guest Lecture: IB Interconnects/PascalBB
  - SI 2009 Guest Lecture: IB Interconnects/PascalBB
  - SI 2010 Guest Lecture: IB Interconnects/PascalBB
- Andrew Shewmaker HPC-5
  - SI 2010 Guest Lecture
- Josip Loncaric, HPC-5
  - SI 2008 Guest Lecture: RR System Management, Prog. Complexity, power efficiency and why it matters
  - SI 2009 Guest Lecture: RR System Management, Prog. Complexity, power efficiency and why it matters
  - SI 2010 Guest Lecture: Cielo: Next Generation Capability Computing Platform for NNSA/ASC
- John Daly HPC-5
  - SI 2008 Guest Lecture
- Parks Fields, HPC-5
  - SI 2008 Guest Lecture: IB Interconnects/PascalBB
  - SI 2009 Guest Lecture: IB Interconnects/PascalBB
  - SI 2010 Guest Lecture: IB Interconnects/PascalBB
- Gary Grider, HPC-DO
  - SI 2007 Guest Lecture: Summer Institute Overview
  - SI 2008 Guest Lecture: Summer Institute Overview
  - SI 2009 Guest Lecture: Summer Institute Overview
  - SI 2010 Guest Lecture: Summer Institute Overview
- Manuel Vigil, HPC-DO
  - SI 2008 Guest Lecture: History of Super Computing
  - SI 2009 Guest Lecture: History of Super Computing
  - SI 2010 Guest Lecture: History of Super Computing

Cluster and Network Summer Institute Project Mentor(s)

- Alfred Torrez, HPC-1
  - SI2010 Student Team Project: Performance Analysis and Evaluation of LANL's PaScalBB IO nodes using QDR Infiniband and Multiple 10-Gigabit Ethernets
- Susan Coulter, HPC-3
  - SI2007 Student Team Project: Advanced Health Monitoring of Computer Clusters
  - SI2010 Student Team Project: Implementation & Comparison of RDMA Over Ethernet
- David Kratzer, HPC-3
  - SI2007 Student Team Project: Xen Virtualization in an Enterprise Environment
- Kevin Tegtmeier, HPC-3
  - SI2008 Student Team Project: High Performance Computing with Configuration Management
- Ben McClelland, HPC-5
  - SI2009 Student Team Project: Kerberized NFS
  - SI2009 Student Team Project: Booting Over Infiniband with Perceus Cluster Management
  - SI2010 Student Team Project: Building a private cloud with OpenNebula
  - SI2010 Student Team Project: Implementation & Comparison of RDMA Over Ethernet
- Jharrod LaFon, HPC-3
  - SI2010 Student Team Project: Implementation & Comparison of RDMA Over Ethernet
- Dane Gardner, HPC-3
  - SI2010 Student Team Project: On-Demand Scientific Computing
- Graham Van Heule, HPC-3
  - SI2010 Student Team Project: On-Demand Scientific Computing
- HB Chen, HPC-5
  - SI2007 Student Team Project: Comparison of 10GigE and Infiniband Interconnects
  - SI2007 Student Team Project: Virtualization in an Enterprise Environment
  - SI2008 Student Team Project: High Speed Interconnection Network Performance Studies
  - SI2009 Student Team Project: QDR Infiniband Network Study
  - SI2010 Student Team Project: Performance Analysis and Evaluation of LANL's PaScalIBB IO nodes using QDR Inifiniband and Multiple 10-Gigabit Ethernets
- Andrew Shewmaker, HPC-5
  - SI2009 Student Team Project: Multiprocessor Affinity
  - SI2009 Student Team Project: Booting Over Infiniband with Perceus Cluster Management
  - SI2010 Student Team Project: On-Demand Scientific Computing
- Parks Fields, HPC-5
  - SI2009 Student Team Project: QDR Infiniband Network Study
  - SI2010 Student Team Project: Performance Analysis and Evaluation of LANL's PaScalIBB IO nodes using QDR Inifiniband and Multiple 10-Gigabit Ethernets

Mentor Opportunity

- Mentor of Student for dissertation: Ringer: Distributed Naming on a Global Scale (Shelly Spearing, HPC-1; Jorge Roman, HPC-1)
- Mentor of Student for dissertation: Proper Plugin Protocols (David R. Montoya, HPC-3; Steve Painter, HPC-1; Scott Matthews, HPC-1)
- Mentor of Student for dissertation: Investigating Efficient Real-time Performance Guarantees on Storage Networks (Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- Mentor of Student for dissertation: Data Reliability Techniques for Specialized Storage Environments (Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- Mentor of Student for dissertation: On-line Index Selection for Physical Database Tuning (Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- Mentor of Student for dissertation: Efficient Performance Guarantees on Storage Networks (Gary Grider, HPC-DO; James Nuñez, HPC-5; John Bent, HPC-5)
- Mentor of Student for dissertation: Designing Physical Primitives For Secure Communication In Wireless Sensor Networks (Gary Grider, HPC-DO; Carolyn Connor, HPC-5)
- Mentor of Student for dissertation: Xquery Over Scientific Data Formats (Gary Grider, HPC-DO; John Bent, HPC-5; Carolyn Connor, HPC-5)

Seminars/presentations/lectures

- David Gunter, HPC-3 Presentation: Parallel Tools Talk to ISTI Summer School, July 2008
- Kevin Tegtmeier, HPC-3
  - Presentation: Cluster and Configuration Mgmt for ISTI Summer School, July 2008
  - Presentation: ISTI Cluster Network Summer School Diskless to HPC/CTN, August 2008
- Ben McClelland, HPC-3
  - Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Netbooting Large Clusters Over Infiniband, August 2009
  - Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
  - Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
  - Presentation: CMU PDL Visit Day—PLFS: A Checkpoint Filesystem for Parallel Applications, May 2010
  - Presentation: 2010 Mini-Showcase—Implementation & Comparison of RDMA Over Ethernet, August 2010
- Jharrod LaFon, HPC-3 Presentation: 2010 Mini-Showcase—Implementation & Comparison of RDMA Over Ethernet, August 2010
- Dane Gardner, HPC-3 Presentation: 2010 Mini-Showcase— On-Demand Scientific Computing, August 2010
- Graham Van Heule, HPC-3 Presentation: 2010 Mini-Showcase— On-Demand Scientific Computing, August 2010
- David R. Montoya, HPC-3 Presentation: SPLASH/OOPSLA 2010: Verifying Configuration Files, October 2010

◊ James Nuñez, HPC-5
- Paper: SC09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- Paper: Data Reliability Techniques for Specialized Storage Environments, March 2009
- Paper: IEEE Cluster 2009: Overlapped Checkpointing with Hardware Assist, September 2009
- Presentation: SIAM08 High End Computing File Systems and I/O (HEC FSIO): Coordinating the US Government Research Investments, 2008
- Presentation: ISW08 IO Scalable Forwarding Layer, April 2008
- Presentation: ISW07 Reliability at Scale, no month 2007
- Presentation: SC07-PDSW New Failure Data Releases, no month 2007
- Presentation: CMU PDL Day PLFS: A Checkpoint Filesystem for Parallel Applications, May 2009
- Presentation: CMU PDL Visit Day—PLFS: A Checkpoint Filesystem for Parallel Applications, May 2010
- Presentation: CMU PDL Day PLFS 2009: A Stackable, Log Structured Filesystem for Checkpointing, May 2009
- Presentation: USENIX FAST 09: Overlapped HPC Checkpointing with Hardware Assist, February 2009
- Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Managing High-Bandwidth Real-Time Streaming Data, August 2009
- Presentation: Data Intensive Supercomputing HPC Application Exploration, August 2009
- Presentation: ISSDM Day: Improving RAID-Based Storage Systems with Flash Memory, October 2009
- Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
- Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
- Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
- Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010

◊ John Bent, HPC-5
- Paper: Introducing Map-Reduce to High End Computing PDSW08, November 2008
- Paper: SC09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- Paper: Data Reliability Techniques for Specialized Storage Environments, March 2009
- Paper: Building a parallel file system simulator, June 2009
- Paper: ...And eat it too: High read performance in write-optimized HPC I/O middleware file formats, November 2009
- Paper: Comparing the Performance of Different Parallel Filesystem Placement Strategies, February 2009
- Paper: Building a Parallel file System Simulator, August 2009
- Paper: Mixing Hadoop and HPC Workloads on Parallel Filesystems, November 2009
- Paper: Mahanaxar: Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010
- Presentation: IEEE MSST07 LANL and HEC FSIO Activities, no month 2008
- Presentation: PDL07 Metadata-indexed Cluster Filesystems, no month 2007
- Presentation: PDL07 Metadata-indexed Cluster Filesystems, no month 2007
- Presentation: Introducing Map-Reduce to High End Computing PDSW 08, November 2008
- Presentation: Early Experiences with Disc-style Computation PDL08, October 2008
- Presentation: Early Experiences with Disc-style Computation PDSW08, October 2008
- Presentation: FAST 09 WiP: File System Load Balancing, February 2009
- Presentation: CMU PDL Day PLFS: A Checkpoint Filesystem for Parallel Applications, May 2009
- Presentation: CMU PDL Day PLFS 2009: A Stackable, Log Structured Filesystem for Checkpointing, May 2009
- Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Managing High-Bandwidth Real-Time Streaming Data, August 2009
- Presentation: Multiphase Plan for the IS&T Data Exploratory, August 2009
- Presentation: Data Intensive Supercomputing HPC Application Exploration, August 2009
- Presentation: ISSDM Day: Improving RAID-Based Storage Systems with Flash Memory, October 2009
- Presentation: ISSDM Day: Performance Analysis of Mixed Distributed Filesystem Workloads, October 2009
- Presentation: ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2009
- Presentation: ISSDM Day: PLFS: Parallel LFS, October 2009
- Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
- Presentation: PDL09: PLFS and HDFS: Enabling Parallel Filesystems Semantics in the Cloud, November 2009
- Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
- Presentation: EuroSys 2010—RAID4S: Adding SSDs to RAID Arrays, April 2010
- Presentation: USENIX FAST 10—RAID4S: Adding SSDs to RAID Arrays, February 2010
- Presentation: 26th IEEE Symposium (MSST2010): Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010
- Presentation: CMU PDL Visit Day—PLFS: A Checkpoint Filesystem for Parallel Applications, May 2010
- Presentation: CMU PDL Visit Day—PLFS and HDFS: Enabling Parallel Filesystem Semantics In The Cloud, May 2010
- Presentation: EuroSys 2010—RAID4S: Adding SSDs to RAID Arrays, April 2010
- Presentation: USENIX FAST 10—Enabling Scientific Application I/O on Cloud FileSystems, February 2010
- Presentation: USENIX FAST 10—Enabling Scientific Application I/O on Cloud FileSystems, February 2010
- Presentation: 2010 Mini-Showcase—Managing High-Bandwidth Real-Time Streaming Data, August 2010
- Presentation: 2010 Mini-Showcase—Scalable Simulation of Parallel Filesystems, August 2010
- Presentation: 2010 ISSDM Day: On-line Index Selection for Physical Database Tuning, October 2010
- Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
- Presentation: 2010 ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2010
- Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
- Presentation: 2010 ISSDM Day: File System Trace and Replay, October 2010
- Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
- Presentation: 2010 ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2010
- Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
- Presentation: 2010 ISSDM Day: On-line Index Selection for Physical Database Tuning, October 2010

- HB Chen, HPC-5
  - Paper: Mahanaxar: Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010
  - Presentation: HPC Networks and Interconnects to ISTI Summer School, July 2008
  - Presentation: ISTI Cluster Network Summer School Network to CTN/ HPC, July 2008
  - Presentation: HPC/ISTI Summer Student Project Mini-Showcase: QDR Infiniband Network Study, August 2009
  - Presentation: ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2009
  - Presentation: 26th IEEE Symposium (MSST2010): Quality of Service Guarantees in High-Bandwidth, Real-Time Streaming Data Storage, May 2010
  - Presentation: 2010 Mini-Showcase—Performance Analysis and Evaluation of LANL's PaScalBB IO nodes using PCIe Gen 2.0 Quad-Data-Rate Infiniband and Multiple 10-Gigabit Ethernets, August 2010
  - Presentation: 2010 Mini-Showcase—Performance Analysis and Evaluation of LANL's PaScalBB IO nodes using PCIe Gen 2.0 Quad-Data-Rate Infiniband and Multiple 10-Gigabit Ethernets, August 2010
  - Presentation: 2010 Mini-Showcase—Managing High-Bandwidth Real-Time Streaming Data, August 2010
  - Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
  - Presentation: 2010 ISSDM Day: Managing High-Bandwidth Real-Time Data Storage, October 2010
- Andrew Shewmaker, HPC-5
  - Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Processor Affinity, August 2009
  - Presentation: 2010 Mini-Showcase—On-Demand Scientific Computing, August 2010
  - Presentation: 2010 ISSDM Day: Investigating Efficient Real-Time Performance Guarantees on Storage Networks, October 2010
- Carolyn Connor, HPC-5
  - Presentation: Multiphase Plan for the IS&T Data Exploratory, August 2009
  - Presentation: 2010 ISSDM Day: File System Trace and Replay, October 2010
- John Daly, HPC-5
  - Presentation: Scheduling and Resilience to ISTI Summer School, July 2008
  - Parks Fields, HPC-5
  - Presentation: HPC Networks and Interconnects to ISTI Summer School, July 2008
  - Presentation: 2010 Mini-Showcase—Performance Analysis and Evaluation of LANL's PaScalBB IO nodes using PCIe Gen 2.0 Quad-Data-Rate Infiniband and Multiple 10-Gigabit Ethernets, August 2010
- Meghan Wingate, HPC-5
  - Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
  - Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
  - Presentation: CMU PDL Visit Day—PLFS: A Checkpoint Filesystem for Parallel Applications, May 2010
  - Presentation: USENIX FAST 10—Enabling Scientific Application I/O on Cloud FileSystems, February 2010
  - Presentation: 2010 Mini-Showcase—Improving RAID Performance with Solid State Drives, August 2010
  - Presentation: 2010 Mini-Showcase—File System Trace and Replay Redux, August 2010
  - Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
  - Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
  - Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
- Grant Mackey, HPC-5
  - Presentation: Data Intensive Supercomputing HPC Application Exploration, August 2009
- Christopher Mitchell, HPC-5
  - Paper: IEEE Cluster 2009: Overlapped Checkpointing with Hardware Assist, September 2009
  - Presentation: USENIX FAST 09: Overlapped HPC Checkpointing with Hardware Assist , February 2009
  - Presentation: Data Intensive Supercomputing HPC Application Exploration, August 2009
- Gary Grider, HPC-DO
  - Paper: SC09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
  - Paper: Data Reliability Techniques for Specialized Storage Environments, March 2009
  - Presentation: SIAM08 High End Computing File Systems and I/O (HEC FSIO): Coordinating the US Government Research Investments , no month 2008
  - Presentation: SSRC08 - Key Note HECFSIO Status, May 2008
  - Presentation: ISW08 Storage Technology Panel, April 2008
  - Presentation: SC07 Petascale Storage Panel, no month 2007
  - Presentation: PDSW08 Rewarding the Public Release of Valuable Data and Resources, November 2008
  - Presentation: Exa-Yotta Storage SC08, November 2008
  - Presentation: CMU PDL Day PLFS: A Checkpoint Filesystem for Parallel Applications, May 2009
  - Presentation: CMU PDL Day PLFS 2009: A Stackable, Log Structured Filesystem for Checkpointing, May 2009
  - Presentation: HPC/ISTI Summer Student Project Mini-Showcase: Managing High-Bandwidth Real-Time Streaming Data, August 2009
  - Presentation: Multiphase Plan for the IS&T Data Exploratory, August 2009
  - Presentation: Data Intensive Supercomputing HPC Application Exploration, August 2009
  - Presentation: ISSDM Day: Improving RAID-Based Storage Systems with Flash Memory, October 2009
  - Presentation: PDL Workshop & Retreat PLFS: A Checkpoint Filesystem for Parallel Applications, October 2009
  - Presentation: PDL09: PLFS: A Checkpoint Filesystem for Parallel Applications, November 2009
  - Presentation: CMU PDL Visit Day—PLFS: A Checkpoint Filesystem for Parallel Applications, May 2010
  - Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
  - Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
  - Presentation: 2010 ISSDM Day: File System Trace and Replay, October 2010
  - Presentation: 2010 ISSDM Day: RAID4S Hardware Performance with a Linux Software RAID, October 2010
  - Presentation: 2010 ISSDM Day: Scalable Simulation of Parallel File Systems, October 2010
  - Presentation: 2010 ISSDM Day: ExaScale
  - FSIO - Can we get there? Can we afford to?, October 2010
- Manuel Vigil, HPC-DO
  - Presentation: Supercomputing at LANL for ISTI Summer School, July 2008

UCSC Graduate Classes are offered through ISTI (for credit toward a graduate degree in computer science or for professional development):

- Alfred Torrez, HPC-1: Attended course: Storage Systems, winter 2007
- Douglas Coombs, HPC-1: Attended course: Database Systems I: Principles of Database Systems, fall 2008
- Wayne Scoggins, HPC-1: Attended course: Machine Learning, winter 2009
- David Montoya, HPC-3: Attended course: Computer Architecture, fall 2008

- James Nuñez, HPC-5
  - Attended course: Analysis of Algorithms, fall 2006
  - Attended course: Advanced Topics in Computer Systems: Activation Object-Based Storage, fall 2008
- Andrew Shewmaker, HPC-5
  - Attended course: Analysis of Algorithms, fall 2006
  - Attended course: Advanced Topics in Computer Systems: Activation Object-Based Storage, fall 2008
  - Attended course: Computer Architecture, winter 2007
  - Attended course: Distributed Systems, spring 2007
  - Attended course: Advanced Operating Systems, fall 2006
  - Attended course: Seminar in Computer Systems, fall 2009
  - Attended course: Seminar in Computer Systems, winter 2010
  - Attended course: Seminar in Computer Systems, spring 2010
  - Attended course: Seminar in Computer Systems, fall 2010
- Daryl Grunau, HPC-5: Attended course: Analysis of Algorithms, fall 2006
- Latchesar Ionkov, HPC-5
  - Attended course: Analysis of Algorithms, fall 2006
  - Attended course: Distributed Systems, spring 2007
  - Attended course: Principles of Database Systems, fall 2007
  - Attended course: Parallel Processing, winter 2008
  - Attended course: Advanced Operating Systems, fall 2006
  - Attended course: Computer Architecture, fall 2008
  - Attended course: Computer Networks, fall 2008
  - Attended course: Programming Languages and Environments, winter 2009
  - Attended course: Sensor Networks, winter 2009
  - Attended course: Coalescing Analysis & Storage, fall 2010
- David Dubois, HPC-5
  - Attended course: Analysis of Algorithms, fall 2006
  - Attended course: Independent Study Storage systems, winter 2007
- Sean Blanchard, HPC-5
  - Attended course: Advanced Operating Systems, fall 2006
  - Attended course: Coalescing Analysis & Storage, fall 2010
- Satsangat Khalsa, HPC-5
  - Attended course: Independent Study Storage systems, winter 2007
  - Attended course: Storage Systems, winter 2007
- Tony Heaton, HPC-5: Attended course: Advanced Operating Systems, fall 2006
- Hugh Greenberg, HPC-5
  - Attended course: Advanced Computer Security, spring 2009
  - Attended course: Coalescing Analysis & Storage, fall 2010